


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search:  The ACM Digital Library  The Guide


[THE ACM DIGITAL LIBRARY](#)

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

 Terms used [HTML AND condensing AND optimization](#)

Found 5,203 of 132,857

Sort results by

 [Save results to a Binder](#)
[Try an Advanced Search](#)

Display results

 [Search Tips](#)
[Try this search in The ACM Guide](#)
 [Open results in a new window](#)
(4/29/2004) 09/683,042

Results 1 - 20 of 200

 Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

      Relevance scale

## 1 [GALAHAD, a library of thread-safe Fortran 90 packages for large-scale nonlinear optimization](#)

Nicholas I. M. Gould, Dominique Orban, Philippe L. Toint

 December 2003 **ACM Transactions on Mathematical Software (TOMS)**, Volume 29 Issue 4

 Full text available: [pdf\(146.51 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We describe the design of version 1.0 of GALAHAD, a library of Fortran 90 packages for large-scale nonlinear optimization. The library particularly addresses quadratic programming problems, containing both interior point and active set algorithms, as well as tools for preprocessing problems prior to solution. It also contains an updated version of the venerable nonlinear programming package, LANCELOT.

**Keywords:** Fortran 90, GALAHAD, LANCELOT, large-scale nonlinear optimization, large-scale quadratic programming

## 2 [Face recognition: A literature survey](#)

W. Zhao, R. Chellappa, P. J. Phillips, A. Rosenfeld

 December 2003 **ACM Computing Surveys (CSUR)**, Volume 35 Issue 4

 Full text available: [pdf\(4.28 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

As one of the most successful applications of image analysis and understanding, face recognition has recently received significant attention, especially during the past several years. At least two reasons account for this trend: the first is the wide range of commercial and law enforcement applications, and the second is the availability of feasible technologies after 30 years of research. Even though current machine recognition systems have reached a certain level of maturity, their success is ...

**Keywords:** Face recognition, person identification

## 3 [Document formatting: Creating personalized documents: an optimization approach](#)

Lisa Purvis, Steven Harrington, Barry O'Sullivan, Eugene C. Freuder

 November 2003 **Proceedings of the 2003 ACM symposium on Document engineering**

 Full text available: [pdf\(432.32 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The digital networked world is enabling and requiring a new emphasis on personalized document creation. The new, more dynamic digital environment demands tools that can reproduce both the contents and the layout automatically, tailored to personal needs and transformed for the presentation device, and can enable novices to easily create such documents. In order to achieve such automated document assembly and transformation, we have formalized custom document creation as a multiobjective optimiza ...


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
 The ACM Digital Library  The Guide


[THE ACM DIGITAL LIBRARY](#)

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
**Terms used** [HTML AND compressing strings](#)

Found 5,460 of 132,857

Sort results by

 Save results to a Binder

Display results

 Search Tips

 Open results in a new window

[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Results 161 - 180 of 200

 Result page: [previous](#) [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

### **161 Improving end-to-end performance of the Web using server volumes and proxy filters**

Edith Cohen, Balachander Krishnamurthy, Jennifer Rexford

 October 1998 **ACM SIGCOMM Computer Communication Review , Proceedings of the ACM SIGCOMM '98 conference on Applications, technologies, architectures, and protocols for computer communication**, Volume 28 Issue 4

 Full text available: [pdf\(1.79 MB\)](#)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The rapid growth of the World Wide Web has caused serious performance degradation on the Internet. This paper offers an end-to-end approach to improving Web performance by collectively examining the Web components --- clients, proxies, servers, and the network. Our goal is to reduce user-perceived latency and the number of TCP connections, improve cache coherency and cache replacement, and enable prefetching of resources that are likely to be accessed in the near future. In our scheme, server re ...

**Keywords:** Web, caching, coherency, filters, piggybacking, prefetching, volumes

### **162 Managing the software design documents with XML**

Junichi Suzuki, Yoshikazu Yamamoto

 September 1998 **Proceedings of the 16th annual international conference on Computer documentation**

 Full text available: [pdf\(1.09 MB\)](#)

 Additional Information: [full citation](#), [references](#), [index terms](#)
**Keywords:** CASE data interchange, UML, XML, software model interchange

### **163 Faster IP lookups using controlled prefix expansion**

V. Srinivasan, George Varghese

 June 1998 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1998 ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems**, Volume 26 Issue 1

 Full text available: [pdf\(1.31 MB\)](#)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Internet (IP) address lookup is a major bottleneck in high performance routers. IP address lookup is challenging because it requires a *longest matching prefix* lookup. It is compounded by increasing routing table sizes, increased traffic, higher speed links, and the migration to 128 bit IPv6 addresses. We describe how IP lookups can be made faster using a new technique called *controlled prefix expansion*. Controlled prefix expansion, together with optimization techniques based on dyn ...

**Keywords:** automated layout, constrained optimization, constraint-based reasoning, document design, genetic algorithm, multiobjective optimization

4 Optimizing document format: Two diet plans for fat PDF

Thomas A. Phelps, Robert Wilensky

November 2003 **Proceedings of the 2003 ACM symposium on Document engineering**

Full text available:  pdf(198.98 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

As Adobe's Portable Document Format has exploded in popularity so too has the number of PDF generators, and predictably the quality of generated PDF varies considerably. This paper surveys a range of PDF optimizations for space, and reports the results of a tool that can postprocess existing PDFs to reduce file sizes by 20 to 70% for large classes of PDFs. (Further reduction can often be obtained by recoding images to lower resolutions or with newer compression methods such as JBIG2 or JPEG2000, bu ...

**Keywords:** PDF, compact PDF, compression, multivalent

5 Image annotation and video summarization: Generation of interactive multi-level video summaries

Frank Shipman, Andreas Grgensohn, Lynn Wilcox

November 2003 **Proceedings of the eleventh ACM international conference on Multimedia**

Full text available:  pdf(522.78 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper, we describe how a detail-on-demand representation for interactive video is used in video summarization. Our approach automatically generates a hypervideo composed of multiple video summary levels and navigational links between these summaries and the original video. Viewers may interactively select the amount of detail they see, access more detailed summaries, and navigate to the source video through the summary. We created a representation for interactive video that supports a wi ...

**Keywords:** hypervideo, link generation, video editing, video summarization

6 GADGET: a toolkit for optimization-based approaches to interface and display generation

James Fogarty, Scott E. Hudson

November 2003 **Proceedings of the 16th annual ACM symposium on User interface software and technology**

Full text available:  pdf(823.58 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Recent work is beginning to reveal the potential of numerical optimization as an approach to generating interfaces and displays. Optimization-based approaches can often allow a mix of independent goals and constraints to be blended in ways that would be difficult to describe algorithmically. While optimization-based techniques appear to offer several potential advantages, further research in this area is hampered by the lack of appropriate tools. This paper presents GADGET, an experimental toolkit ...

**Keywords:** display generation, layout algorithms, numerical optimization, perceptually optimized displays, toolkits

7 Extending Java for high-level Web service construction

Aske Simon Christensen, Anders Møller, Michael I. Schwartzbach

November 2003 **ACM Transactions on Programming Languages and Systems (TOPLAS)**,

Volume 25 Issue 6

Full text available:  pdf(947.02 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We incorporate innovations from the <bigwig> project into the Java language to provide high-level features for Web service programming. The resulting language, JWIG, contains an advanced session model and a flexible mechanism for dynamic construction of XML documents, in particular XHTML. To support program development we provide a suite of program analyses that at compile time verify for a given program that no runtime errors can occur while building documents or receiving form input, and ...

**Keywords:** Interactive Web services, XML, data-flow analysis

**8 Compiler optimizations for power, performance: Tracking object life cycle for leakage energy optimization**

G. Chen, N. Vijaykrishnan, M. Kandemir, M. J. Irwin, M. Wolczko

October 2003 **Proceedings of the 1st IEEE/ACM/IFIP international conference on Hardware/software codesign & system synthesis**

Full text available: [pdf\(304.66 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The focus of this work is on utilizing the state of objects during their lifespan in optimizing the leakage energy consumed in the data caches when executing embedded Java applications. Our analysis reveals that a major portion of the leakage energy is actually wasted in retaining the objects beyond their last use. In order to eliminate this wastage, we investigate three approaches that use the garbage collector, escape analysis and last use analysis for reducing leakage energy. Finally, we trac ...

**Keywords:** Java, cache, leakage energy

**9 Modeling and validation of service-oriented architectures: application vs. style**

Luciano Baresi, Reiko Heckel, Sebastian Thöne, Dániel Varró

September 2003 **ACM SIGSOFT Software Engineering Notes , Proceedings of the 9th European software engineering conference held jointly with 10th ACM SIGSOFT international symposium on Foundations of software engineering**, Volume 28 Issue 5

Full text available: [pdf\(239.74 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Most applications developed today rely on a given *middleware platform* which governs the interaction between components, the access to resources, etc. To decide, which platform is suitable for a given application (or more generally, to understand the interaction between application and platform), we propose UML models of both the *architectural style of the platform* and the application scenario. Based on a formal interpretation of these as graphs and *graph transformation systems*

**Keywords:** architectural style, graph transformation system, model checking, modeling middleware platforms, service-oriented architectures

**10 Metaheuristics in combinatorial optimization: Overview and conceptual comparison**

Christian Blum, Andrea Roli

September 2003 **ACM Computing Surveys (CSUR)**, Volume 35 Issue 3

Full text available: [pdf\(431.84 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The field of metaheuristics for the application to combinatorial optimization problems is a rapidly growing field of research. This is due to the importance of combinatorial optimization problems for the scientific as well as the industrial world. We give a survey of the nowadays most important metaheuristics from a conceptual point of view. We outline the different components and concepts that are used in the different metaheuristics in order to analyze their similarities and differences. Two v ...

**Keywords:** Metaheuristics, combinatorial optimization, diversification., intensification

**11 Research track: Classifying large data sets using SVMs with hierarchical clusters**

Hwanjo Yu, Jiong Yang, Jiawei Han

August 2003 **Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining**Full text available: [pdf\(381.15 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Support vector machines (SVMs) have been promising methods for classification and regression analysis because of their solid mathematical foundations which convey several salient properties that other methods hardly provide. However, despite the prominent properties of SVMs, they are not as favored for large-scale data mining as for pattern recognition or machine learning because the training complexity of SVMs is highly dependent on the size of a data set. Many real-world data mining applicati ...

**Keywords:** hierarchical cluster, support vector machines

**12 Memory-based and disk-based algorithms for very high degree permutation groups**

Gene Cooperman, Eric Robinson

August 2003 **Proceedings of the 2003 international symposium on Symbolic and algebraic computation**Full text available: [pdf\(222.98 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Group membership is a fundamental algorithm, upon which most other algorithms of computational group theory depend. Until now, group membership for permutation groups has been limited to ten million points or less. We extend the applicability of group membership algorithms to permutation groups acting on more than 100,000,000 points. As an example, we experimentally construct a group membership data structure for Thompson's group, acting on 143,127,000 points, in 36 minutes. More significantly, ...

**Keywords:** Thompson's group, disk-based methods, group membership, permutation groups, permutation multiplication

**13 Speculative execution: A new speculation technique to optimize floating-point performance while preserving bit-by-bit reproducibility**

Mikio Takeuchi, Hideaki Komatsu, Toshio Nakatani

June 2003 **Proceedings of the 17th annual international conference on Supercomputing**Full text available: [pdf\(227.01 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The bit-by-bit reproducibility of floating-point results, which is defined by the IEEE 754 standard, prohibits optimizations such as reassociation and the use of native operations such as fused multiply-add (FMA), and thus it significantly impairs floating-point performance. Recent network-oriented languages such as Java strictly conform to the standard, and thus their numerical computing performance becomes inherently lower than conventional languages. In this paper, we propose a new software te ...

**Keywords:** IA-64, IEEE 754, Java, accuracy, bit-by-bit reproducibility, floating-point speculation, fused multiply-add, instruction-level parallelism, just-in-time compiler, loop unrolling, prefetching, privatization, reassociation, software pipelining, striding

**14 Spatial and nearest-neighbor queries: Contorting high dimensional data for efficient main memory KNN processing**

Bin Cui, Beng Chin Ooi, Jianwen Su, Kian-Lee Tan

June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on Management of data**Full text available: [pdf\(229.89 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper, we present a novel index structure, called  $\Delta$ -tree, to speed up processing of high-dimensional K-nearest neighbor (KNN) queries in main memory environment. The  $\Delta$ -

tree is a multi-level structure where each level represents the data space at different dimensionalities: the number of dimensions increases towards the leaf level which contains the data at their full dimensions. The remaining dimensions are obtained using *Principal Component Analysis*, which has the des ...

### **15 Session 11A: Extractors: optimal up to constant factors**

Chi-Jen Lu, Omer Reingold, Salil Vadhan, Avi Wigderson

June 2003 **Proceedings of the thirty-fifth ACM symposium on Theory of computing**

Full text available:  pdf(246.08 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper provides the first explicit construction of extractors which are simultaneously optimal up to constant factors in *both* seed length and output length. More precisely, for every  $n, k$ , our extractor uses a random seed of length  $O(\log n)$  to transform any random source on  $n$  bits with (min-)entropy  $k$ , into a distribution on  $(1-\alpha)k$  bits that is  $\epsilon$ -close to uniform. Here  $\alpha$  and  $\epsilon$  can be taken to be any positive  $c$  ...

**Keywords:** condensers, locally decodable error-correcting codes, mergers, pseudorandomness, randomness extractors

### **16 Optimization: Optimized color gamuts for tiled displays**

Marshall Bern, David Eppstein

June 2003 **Proceedings of the nineteenth conference on Computational geometry**

Full text available:  pdf(155.36 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We consider the problem of finding a large color space that can be generated by all units in multi-projector tiled display systems. Viewing the problem geometrically as one of finding a large parallelepiped within the intersection of multiple parallelepipeds, and using colorimetric principles to define a volume-based objective function for comparing feasible solutions, we develop an algorithm for finding the optimal gamut in time  $O(n^3)$ , where  $n$  denotes the number of proj ...

**Keywords:** additive color, color gamuts, gamut mapping, geometric optimization, high-resolution display systems, quasiconvex programming, tiled displays

### **17 Reading patterns and usability in visualizations of electronic documents**

Kasper Hornbæk, Erik Frøkjær

June 2003 **ACM Transactions on Computer-Human Interaction (TOCHI)**, Volume 10 Issue 2

Full text available:  pdf(4.18 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present an exploration of reading patterns and usability in visualizations of electronic documents. Twenty subjects wrote essays and answered questions about scientific documents using an overview+detail, a fisheye, and a linear interface. We study reading patterns by progression maps that visualize the progression of subjects' reading activity, and by visibility maps that show for how long different parts of the document are visible. The reading patterns help explain differences in usability ...

**Keywords:** Electronic documents, digital documents, fisheye interface, information retrieval, information visualization, overview+detail interface, reading, reading patterns

### **18 Web crawling and measurement: Efficient URL caching for world wide web crawling**

Andrei Z. Broder, Marc Najork, Janet L. Wiener

May 2003 **Proceedings of the twelfth international conference on World Wide Web**

Full text available:  pdf(174.37 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Crawling the web is deceptively simple: the basic algorithm is (a) Fetch a page (b) Parse it to extract all linked URLs (c) For all the URLs not seen before, repeat (a)-(c). However, the

size of the web (estimated at over 4 billion pages) and its rate of change (estimated at 7% per week) move this plan from a trivial programming exercise to a serious algorithmic and system design challenge. Indeed, these two factors alone imply that for a reasonably fresh and complete crawl of the web, step (a) ...

**Keywords:** URL caching, caching, crawling, distributed crawlers, web crawlers, web graph models

**19 Automatically proving the correctness of compiler optimizations**

Sorin Lerner, Todd Millstein, Craig Chambers

May 2003 **ACM SIGPLAN Notices , Proceedings of the ACM SIGPLAN 2003 conference on Programming language design and implementation**, Volume 38 Issue 5

Full text available: [pdf\(285.83 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We describe a technique for automatically proving compiler optimizations *sound*, meaning that their transformations are always semantics-preserving. We first present a domain-specific language, called Cobalt, for implementing optimizations as guarded rewrite rules. Cobalt optimizations operate over a C-like intermediate representation including unstructured control flow, pointers to local variables and dynamically allocated memory, and recursive procedures. Then we describe a technique for ...

**Keywords:** automated correctness proofs, compiler optimization

**20 Data remapping for design space optimization of embedded memory systems**

Rodric M. Rabbah, Krishna V. Palem

May 2003 **ACM Transactions on Embedded Computing Systems (TECS)**, Volume 2 Issue 2

Full text available: [pdf\(885.05 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this article, we present a novel linear time algorithm for *data remapping*, that is, (i) lightweight; (ii) fully automated; and (iii) applicable in the context of pointer-centric programming languages with dynamic memory allocation support. All previous work in this area lacks one or more of these features. We proceed to demonstrate a *novel application of this algorithm as a key step in optimizing the design of an embedded memory system*. Specifically, we show that by virtue of lo ...

**Keywords:** Design space exploration, caches, compiler optimization, data remapping, embedded systems, memory hierarchy, memory subsystem

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

**164 Herbarium specimen browser: a tool for accessing botanical specimen collections**

Erich R. Schneider, John J. Leggett, Richard K. Furuta, Hugh D. Wilson, Stephan L. Hatch  
 May 1998 **Proceedings of the third ACM conference on Digital libraries**

Full text available:  pdf(1.12 MB) Additional Information: [full citation](#), [references](#), [index terms](#)

**165 Beyond SGML**

Roger Price  
 May 1998 **Proceedings of the third ACM conference on Digital libraries**

Full text available:  pdf(1.12 MB) Additional Information: [full citation](#), [references](#), [index terms](#)

**166 CiteSeer: an automatic citation indexing system**

C. Lee Giles, Kurt D. Bollacker, Steve Lawrence  
 May 1998 **Proceedings of the third ACM conference on Digital libraries**

Full text available:  pdf(1.18 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

**167 Improving data-flow analysis with path profiles**

Glenn Ammons, James R. Larus  
 May 1998 **ACM SIGPLAN Notices , Proceedings of the ACM SIGPLAN 1998 conference on Programming language design and implementation**, Volume 33 Issue 5

Full text available:  pdf(1.57 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Data-flow analysis computes its solutions over the paths in a control-flow graph. These paths---whether feasible or infeasible, heavily or rarely executed---contribute equally to a solution. However, programs execute only a small fraction of their potential paths and, moreover, programs' execution time and cost is concentrated in a far smaller subset of *hot paths*. This paper describes a new approach to analyzing and optimizing programs, which improves the precision of data flow analysis al ...

**168 An efficient text input method for pen-based computers**

Toshiyuki Masui  
 January 1998 **Proceedings of the SIGCHI conference on Human factors in computing systems**

Full text available:  pdf(1.06 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

**Keywords:** POBox, hand-held devices, input devices, international interfaces, pen-based input, predictive interface

**169 Bibliography of recent publications on computer communication**

Martha Steenstrup  
 January 1998 **ACM SIGCOMM Computer Communication Review**, Volume 28 Issue 1

Full text available:  pdf(2.02 MB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

The quantitative results presented in our SIGCOMM '97 paper [1] include numerous minor errors. These errors were caused by programming bugs that led to faulty analyses and simulations, and by inaccurate transcriptions during the preparation of the paper. Here we present corrected figures and tables, as well as corrections to values that appeared in the text of the original paper. The effect of correcting the errors is to reduce the differences between the results based on the proxy trace and tho ...

**170 Fast detection of communication patterns in distributed executions**

Thomas Kunz, Michiel F. H. Seuren

November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research**Full text available:  pdf(4.21 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

**171 An object-oriented SGML/HyTime compliant multimedia database management system**M. Tamer Özsu, Paul Iglinski, Duane Szafron, Sherine El-Medani, Manuela Junghanns  
November 1997 **Proceedings of the fifth ACM international conference on Multimedia**Full text available:  pdf(1.77 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)**172 Cluster-based scalable network services**Armando Fox, Steven D. Gribble, Yatin Chawathe, Eric A. Brewer, Paul Gauthier  
October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles**, Volume 31 Issue 5Full text available:  pdf(2.42 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)**173 Application performance and flexibility on exokernel systems**M. Frans Kaashoek, Dawson R. Engler, Gregory R. Ganger, Héctor M. Briceño, Russell Hunt, David Mazières, Thomas Pinckney, Robert Grimm, John Jannotti, Kenneth Mackenzie  
October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles**, Volume 31 Issue 5Full text available:  pdf(2.39 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)**174 Web based teaching: a minimalist approach**Chris Jesshope  
July 1997 **Proceedings of the second Australasian conference on Computer science education**Full text available:  pdf(1.04 MB) Additional Information: [full citation](#), [references](#), [index terms](#)**175 LDC online: a digital library for linguistic research and development**Zhibiao Wu, Mark Liberman  
July 1997 **Proceedings of the second ACM international conference on Digital libraries**Full text available:  pdf(848.96 KB) Additional Information: [full citation](#), [references](#), [index terms](#)**176 An OpenMath 1.0 implementation**Stéphane Dalmas, Marc Gaétano, Stephen Watt  
July 1997 **Proceedings of the 1997 international symposium on Symbolic and algebraic computation**Full text available:  pdf(1.03 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

**177 Browsing in digital libraries: a phrase-based approach**

Craig G. Nevill-Manning, Ian H. Witten, Gordon W. Paynter

July 1997 **Proceedings of the second ACM international conference on Digital  
libraries**Full text available:  pdf(1.22 MB)Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)**178 The SimpleScalar tool set, version 2.0**

Doug Burger, Todd M. Austin

June 1997 **ACM SIGARCH Computer Architecture News**, Volume 25 Issue 3Full text available:  pdf(985.46 KB)Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

This document describes release 2.0 of the SimpleScalar tool set, a suite of free, publicly available simulation tools that offer both detailed and high-performance simulation of modern microprocessors. The new release offers more tools and capabilities, precompiled binaries, cleaner interfaces, better documentation, easier installation, improved portability, and higher performance. This paper contains a complete description of the tool set, including retrieval and installation instructions, a d ...

**179 Code compression**

Jens Ernst, William Evans, Christopher W. Fraser, Todd A. Proebsting, Steven Lucco

May 1997 **ACM SIGPLAN Notices , Proceedings of the ACM SIGPLAN 1997 conference  
on Programming language design and implementation**, Volume 32 Issue 5Full text available:  pdf(1.11 MB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Current research in compiler optimization counts mainly CPU time and perhaps the first cache level or two. This view has been important but is becoming myopic, at least from a system-wide viewpoint, as the ratio of network and disk speeds to CPU speeds grows exponentially. For example, we have seen the CPU idle for most of the time during paging, so compressing pages can increase total performance even though the CPU must decompress or interpret the page contents. Another profile shows that many ...

**180 Designing Dexter-based hypermedia services for the World Wide Web**

Kaj Grønbæk, Niels Olof Bouvin, Lennert Sloth

April 1997 **Proceedings of the eighth ACM conference on Hypertext**Full text available:  pdf(1.08 MB)Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

**Keywords:** ActiveX, Dexter hypertext reference model, HTML, Java, JavaScript, OLE, World Wide Web, link objects, open hypermedia service

Results 161 - 180 of 200

Result page: [previous](#) [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)